Shorter gate sequences for quantum computing by mixing unitaries

Earl Campbell[1,*]

[1]*Department of Physics and Astronomy, University of Sheffield, Sheffield, UK*

Fault-tolerant quantum computers compose elements of a discrete gate set in order to approximate a target unitary. The problem of minimising the number of gates is known as gate-synthesis. The approximation error is a form of coherent noise, which can be significantly more damaging than comparable incoherent noise. We show how mixing over different gate sequences can convert this coherent noise into an incoherent form. As measured by diamond distance, the post-mixing noise is quadratically smaller than before mixing, with no additional resource cost. Equivalently, we can use a shorter gate sequence to achieve the same precision as unitary gate-synthesis, with a factor $1/2$ reduction for a broad class of problems.

arXiv:1612.02689v2 [quant-ph] 23 Jan 2017

The constraints of fault-tolerant quantum computing mean that the available quantum gates form a discrete set. Such a gate set is said to be universal if it generates a group that gives a dense cover over all unitaries. That is, any target unitary can be approximated to any desired level of precision with a sufficiently long sequence of gates. The Solovay-Kitaev [1–4] theorem ensures that whenever we have a universal gate set, we can achieve a circuit depth that is poly-logarithmic in the inverse precision. The Solovay-Kitaev theorem is a very powerful and general result, but in practice yields very long gate sequences. Remarkable progress beyond Solovay-Kitaev has been made in recent years by focusing on gate-sets that naturally arise in fault-tolerant quantum computing, in particular the Clifford+T gate set, with the flourishing topic becoming known as gate-synthesis [5–8].

A common feature of both new and old approaches to gate-synthesis is the approximation of the target unitary with a different unitary. Then the approximation error is a form of coherent noise, which has attracted attention as being especially pernicious to quantum computations [9, 10]. It has, however, been observed several times that mixing over equivalent circuits can average out coherent noise into less damaging incoherent noise [11–15]. For instance, when the individual gates suffer from coherent noise, randomized compiling has been shown to quadratically reduce this noise source [14]. In the context of gate synthesis, the approximation error appears even when the components of our gate set are perfect, and so a different approach is required.

Here we give the first general set of tools for mixing out the approximation errors in gate synthesis. Quantifying this noise by the diamond norm, we find our approach reduces noise from ϵ to $O(\epsilon^2)$, without increasing circuit depth or any other metric of resource cost. Alternatively, we can achieve $O(\epsilon)$ noise with reduced resource cost. If the resource cost of unitary gate-synthesis scales as $A\log(\epsilon^{-1})^\gamma$, then using quantum channels ϵ noise can be attained with resource cost $A(1/2)^\gamma\log(\epsilon^{-1})^\gamma$ in the small ϵ limit. Many recent gate-synthesis algorithms have $\gamma = 1$ scaling, and so in these setting we cut costs in half. This is an extension of the notion of magic state dilution in Ref. [16], but here applied to synthesis of operations, rather than states. When completing this work, some similar insights were reported by Hastings [17], though without the explicit convex hull finding algorithm provided here.

* earltcampbell@gmail.com

I. NOTATION

We use $||\dots||$ throughout for the operator norm, so that $||X||$ is the largest singular value of X. We also make use of the Schatten 1 norm on operators denoted $||\dots||_1$, which equals the sum of the singular values. Throughout we make use of several norm properties discussed in standard texts [18, 19]. For a quantum channel we use the diamond norm $||\dots||_\diamond$ where

$$||\mathcal{E}||_\diamond := \sup\{||(\mathcal{E}\otimes\mathbb{1})(X)||; ||X||_1 \leq 1\}. \quad (1)$$

The diamond norm induces the diamond distance between two channels $\mathcal{E}$ and $\mathcal{E}'$, so that

$$d_\diamond(\mathcal{E},\mathcal{E}') := \frac{1}{2}||\mathcal{E}-\mathcal{E}'||_\diamond, \quad (2)$$

and is widely used [20] to quantify how well an imperfect channel $\mathcal{E}'$ approximates an ideal, target channel $\mathcal{E}$. The diamond distance is well behaved under composition of channels, allowing it to be used in rigorous proofs, including proofs of the threshold theorem for fault-tolerant quantum computing [21]. Despite the average fidelity gaining popularity and being easily measurable by randomised benchmarking [22–25], various commentators have observed that average fidelity is less meaningful than the diamond distance [9].

In inexact gate synthesis, a sequence of available gates are composed to produce some U that gives a good approximation to a target unitary V. Techniques for gate synthesis typically report the precision of these approximations by taking $U-V$ and evaluating some norm. This prompts us to ask how this notion of precision corresponds to the more versatile diamond distance. Denoting, $\mathcal{U}$ and $\mathcal{V}$ as the channels corresponding to U and V, we have

$$d_\diamond(\mathcal{U},\mathcal{V}) \leq ||U-V||, \quad (3)$$

as shown in Refs. [26, 27]. In general, there is no simple lower bound. For instance, if $U=-V$ then $||U-V||=2$, but $\mathcal{U}=\mathcal{V}$ and so $d_\diamond(\mathcal{U},\mathcal{V})=0$. However, these pathologies only arise when $||U-V||$ is large, and many families of unitaries are well behaved. Consider, for instance, unitaries of the form

$U = e^{i\theta Z}$ and $V = e^{i\theta' Z}$, for small $|\theta - \theta'|$ we find $||U - V||$ is very close to the diamond distance. So while unitary precision and diamond distance are very different measures, they often coincide.

Throughout we will use $\mathcal{G}$ to denote the available gate set, and $\mathfrak{C} : \mathcal{G} \to \mathbb{R}^+$ for the associated cost function. To assess the depth of a circuit we would use a constant cost function $\mathfrak{C}(V) = 1$ for all $V \in \mathcal{G}$. However, for the Clifford+T gate set the T gates can be significantly more expensive than Clifford gates due to the resource overhead of magic state distillation [28–31]. In this setting, one often takes $\mathfrak{C}(T) = 1$ and $\mathfrak{C}(C) = 0$ for all C in the Clifford group. The cost of a gate sequence is then taken to be the numerical sum of the composite gate costs. We also use $\langle \mathcal{G} \rangle$ for the group generated by set $\mathcal{G}$. We say a gate set is finite when $\mathcal{G}$ contains a finite number of elements. Lastly, we will use $\mathrm{Conv}[\ldots]$ to denote the convex hull of a set of operators.

II. RESULTS

Here we present two main results of this paper

Theorem 1 *Let $\mathcal{L}$ be some d dimensional Lie group, which is a subgroup of a unitary group $SU(D)$. Let $\mathcal{G}$ be a finite gate set with cost function $\mathfrak{C} : \mathcal{G} \to \mathbb{R}^+$, such that $\langle \mathcal{G} \rangle$ is a dense cover of $\mathcal{L}$ and $\langle \mathcal{G} \rangle \subset \mathcal{L}$. Assume we have a unitary synthesis algorithm: for every $V \in \mathcal{L}$ and all $\epsilon > 0$ the algorithm outputs a finite sequence $U = W_1 W_2 \ldots W_N \in \langle \mathcal{G} \rangle$, such that*

$$||U - V|| \leq \epsilon, \tag{4}$$

$$\sum_{j=1}^{N} \mathfrak{C}(W_j) \leq f(\epsilon), \tag{5}$$

where f is the worst case cost of the unitary synthesis algorithm. It follows that we can construct a channel of the form

$$\mathcal{E}(\rho) = \sum_{j=1}^{n} p_j U_j \rho U_j^\dagger, \tag{6}$$

where all $U_j \in \langle \mathcal{G} \rangle$ and each have cost upper bounded by $f(\epsilon)$, and provided $\epsilon < 0.01$ the post-mixing noise satisfies

$$d_\diamond(\mathcal{E}, \mathcal{V}) \leq 10\epsilon^2. \tag{7}$$

Therefore, $O(\epsilon^2)$ error is achieved at no extra cost.

The simplest setting is that $\mathcal{L} = SU(D)$, so $d = D$, but we also allow for subgroups with $d < D$. Few gate-synthesis techniques exist for multi-qubit or qudit problems, but our results apply there also. It directly applies to the familiar problem of performing general single-qubit rotations from the Clifford+T gate set. The natural cost function of this gate set is $\mathfrak{C}(T) = 1$ and $\mathfrak{C}(C) = 0$ for all C in the Clifford group. For such a cost function, Ross and Selinger [7] showed that efficient gate synthesis of any single qubit gate is possible with $f_{\mathrm{RS}}(\epsilon) = 9\log_2(\epsilon^{-1}) + O(\log_2(\log_2(\epsilon)))$. Using quantum channels, and no more gates, we can ensure $10\epsilon^2$ precision in diamond distance.

We use the terminology axial rotation for single qubit rotations about the Z axis, and denote the group $\mathcal{L}_{\mathrm{ax}}$. For such rotations the above findings apply with the function f_{RS}. However, the Ross and Selinger algorithm can generate axial rotations at a slightly lower cost with leading order $3\log_2(\epsilon^{-1})$, and other algorithms have been tailored to this special case. So one might anticipate that resource savings could be made by tailoring our approach to axial rotations. We find this is indeed the case, but we cannot blindly apply the above result to algorithms for axial rotations. Note that Thm. 1 does not apply in this setting since the generated group $\langle \mathcal{G} \rangle$ contains gates outside $\mathcal{L}_{\mathrm{ax}}$. That is, with $\mathcal{G}$ as the Clifford+T set, the generated group has gates outside the axial rotation group, so $\langle \mathcal{G} \rangle \not\subset \mathcal{L}_{\mathrm{ax}}$. However, our techniques are straightforwardly extended to such scenarios.

Theorem 2 *Let $\mathcal{L}_{\mathrm{ax}}$ be the group of axial rotations. Let $\mathcal{G}$ be a gate set with cost function $\mathfrak{C} : \mathcal{G} \to \mathbb{R}^+$ with Pauli $Z \in \mathcal{G}$ and $\mathfrak{C}(Z)$=0. Assume we have a unitary synthesis algorithm: for every $V \in \mathcal{L}_{\mathrm{ax}}$ and all $\epsilon > 0$ the algorithm outputs a finite sequence $U = W_1 W_2 \ldots W_n \in \langle \mathcal{G} \rangle$, such that*

$$||U - V|| \leq \epsilon, \tag{8}$$

$$\sum_{j=1}^{N} \mathfrak{C}(W_j) \leq f_{\mathrm{ax}}(\epsilon), \tag{9}$$

where f_{ax} is the worst case cost of the unitary synthesis algorithm. It follows that we can construct a channel of the form

$$\mathcal{E}(\rho) = \sum_{j=1}^{4} p_j U_j \rho U_j^\dagger, \tag{10}$$

where all $U_j \in \langle \mathcal{G} \rangle$ and each have cost upper bounded by $f_{\mathrm{ax}}(\epsilon)$, and provided $\epsilon < 0.01$ the post-mixing noise satisfies

$$d_\diamond(\mathcal{E}, \mathcal{V}) \leq 5\epsilon^2. \tag{11}$$

Therefore, $O(\epsilon^2)$ error is achieved at no extra cost.

This result has a slightly better $5\epsilon^2$ instead of $10\epsilon^2$, but more importantly benefits from using f_{ax} which gives a smaller resource overhead than for general qubit rotations.

Let us reflect on how this free error suppression can be swapped in exchanged for cheaper gate sequences. We instead run our protocol and use gate sequences of cost not exceeding $f(\sqrt{\epsilon/\alpha})$, where α is 5 or 10 depending on which theorem we employ. It follows that the post-mixing noise is bounded by ϵ, but resource costs are reduced. The form of f for unitary gate-synthesis is typically $f(\epsilon) \sim A\log(\epsilon^{-1})^\gamma$ upto a small $O\{\log[\log(\epsilon^{-1})\}$ contribution. Our reduced cost is then

$$\begin{aligned} f(\sqrt{\epsilon/\alpha}) &\sim A\log((\epsilon/\alpha)^{-1/2})^\gamma \\ &\sim A\left(\frac{\log(\epsilon^{-1}) + \log(\alpha)}{2}\right)^\gamma \\ &\sim A\log(\epsilon^{-1})^\gamma \left[\left(\frac{1}{2}\right)\left(1 + \frac{\log(\alpha)}{\log(\epsilon^{-1})}\right)\right]^\gamma. \end{aligned} \tag{12}$$

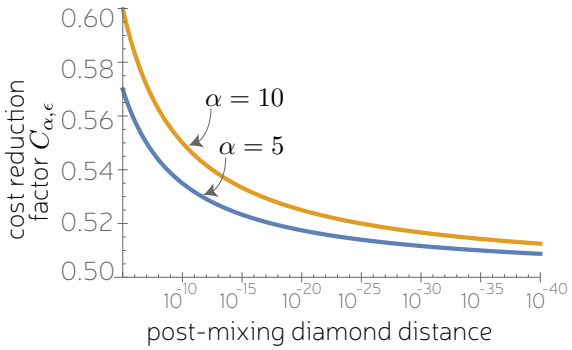

FIG. 1. The resource savings of our approach over unitary gate-synthesis is $C^{\gamma}_{\alpha,\epsilon}$ and here we show $C_{\alpha,\epsilon}$ (see Eq. (13)) for $\alpha = 5, 10$ and a range of post-mixing error rates. The different α correspond to different constant factors in Eq. (7) and Eq. (11).

Therefore, our resource savings are a factor $C^{\gamma}_{\alpha,\epsilon}$ where

$$C_{\alpha,\epsilon} = \left(\frac{1}{2}\right)\left(1 + \frac{\log(\alpha)}{\log(\epsilon^{-1})}\right) \tag{13}$$

collects the terms in the square bracket of Eq. (12). In the small ϵ limit we have $C_{\alpha,\epsilon} \to 1/2$. Typically, ϵ is very small with many algorithms requiring $\epsilon \ll 10^{-10}$ and so $C_{\alpha,\epsilon} \sim 1/2$ is a reasonable approximation. Convergence toward $1/2$ is shown in Fig. 1, with the speed of convergence dictated by α. When proving our theorems we focus on clarity rather than minimising α and believe smaller α is plausible. Lastly, recall that for single qubit problems known algorithms have $\gamma = 1$, but in other settings different γ may appear.

III. THE MIXING LEMMA

Here we prove a Lemma that underpins both Thm. 1 and Thm. 2, and may also enable further extensions.

Lemma 1 *Let V be a target unitary, with associated channel $\mathcal{V}(\rho) = V\rho V^{\dagger}$. Let $a, b > 0$ and $\{U_1, U_2, \ldots, U_n\}$ be a set of unitaries such that*

1. *for all $j \in \{1, \ldots, n\}$ we have $||U_j - V|| \leq a$;*
2. *there exist positive numbers $\{p_j\}$ such that $\sum_{j=1}^{n} p_j = 1$ and $||(\sum_j p_j U_j) - V|| \leq b$.*

It follows that $\mathcal{E} = \sum_j p_j \mathcal{U}_j$ satisfies

$$||\mathcal{E} - \mathcal{V}||_{\diamond} \leq a^2 + 2b. \tag{14}$$

We will find constructions where $a = O(\epsilon)$ and $b = O(\epsilon^2)$, so that the diamond norm is upper bounded by $O(\epsilon^2)$.

For now, we prove the above Lemma. We begin by defining $\delta_j := U_j - V$ so that $||\delta_j|| \leq a$. We also have

$$\sum_j p_j \delta_j = \left(\sum_j p_j U_j\right) - V, \tag{15}$$

with condition (2) of the lemma entailing that $||\sum_j p_j \delta_j|| \leq b$. The channel $\mathcal{E}$ acts as

$$\begin{aligned}\mathcal{E}(X) &= \sum_j p_j U_j X U_j^{\dagger}, \\ &= \sum_j p_j (V + \delta_j) X (V^{\dagger} + \delta_j^{\dagger}).\end{aligned} \tag{16}$$

Since the diamond norm is unitarily invariant, we have $d_{\diamond}(\mathcal{E}, \mathcal{V}) = d_{\diamond}(\mathcal{V}^{\dagger} \circ \mathcal{E}, \mathbb{1})$ where

$$\begin{aligned}(\mathcal{V}^{\dagger} \circ \mathcal{E})(X) &= \sum_j p_j V^{\dagger} U_j X U_j^{\dagger} V \\ &= \sum_j p_j (\mathbb{1} + \tilde{\delta}_j) X (\mathbb{1} + \tilde{\delta}_j^{\dagger}) \\ &= \sum_j p_j (X + \tilde{\delta}_j X + X\tilde{\delta}_j^{\dagger} + \tilde{\delta}_j X \tilde{\delta}_j^{\dagger}),\end{aligned} \tag{17}$$

where $\tilde{\delta}_j := V^{\dagger}\delta_j$. Since the operator norm is unitarily invariant, we have $||\sum_j p_j \tilde{\delta}_j|| = ||\sum_j p_j \delta_j|| \leq b$. Compared to the identity channel $\mathbb{1}$, and using $\sum_j p_j = 1$, we have

$$(\mathcal{V}^{\dagger} \circ \mathcal{E} - \mathbb{1})(X) = \sum_j p_j (\tilde{\delta}_j X + X\tilde{\delta}_j^{\dagger} + \tilde{\delta}_j X \tilde{\delta}_j^{\dagger}). \tag{18}$$

Taking the 1-norm and using the triangle inequality, we have

$$\begin{aligned}||(\mathcal{V}^{\dagger} \circ \mathcal{E} - \mathbb{1})(X)||_1 \leq &||\sum_j p_j \tilde{\delta}_j X||_1 + ||\sum_j p_j X \tilde{\delta}_j^{\dagger}||_1 \\ &+ \sum_j p_j ||\tilde{\delta}_j X \tilde{\delta}_j^{\dagger}||_1.\end{aligned} \tag{19}$$

Using the Hölder inequality and $||X||_1 \leq 1$, we have

$$\begin{aligned}||(\mathcal{V}^{\dagger} \circ \mathcal{E} - \mathbb{1})(X)||_1 \leq &||\sum_j p_j \tilde{\delta}_j|| + ||\sum_j p_j \tilde{\delta}_j^{\dagger}|| \\ &+ \sum_j p_j ||\tilde{\delta}_j|| \cdot ||\tilde{\delta}_j^{\dagger}||.\end{aligned} \tag{20}$$

Noting the property $||M|| = ||M^{\dagger}||$ and condition (1) of Lem. 1, we conclude that $||\tilde{\delta}_j^{\dagger}|| = ||\tilde{\delta}_j|| \leq a$. Therefore, the last sum of terms is upper bounded by a^2. The first two summations are likewise bounded by b by virtue of condition (2). Therefore,

$$||(\mathcal{V}^{\dagger} \circ \mathcal{E} - \mathbb{1})(X)||_1 \leq a^2 + 2b, \tag{21}$$

which is true for all X. If we tensor the channels with the identity this does not affect the proof except to burden the notation, and so

$$||((\mathcal{V}^{\dagger} \circ \mathcal{E} - \mathcal{I}) \otimes \mathcal{I})(X)||_1 \leq a^2 + 2b. \tag{22}$$

Since this is true for all X the diamond norm is also upper bounded by $a^2 + 2b$. This completes the proof.

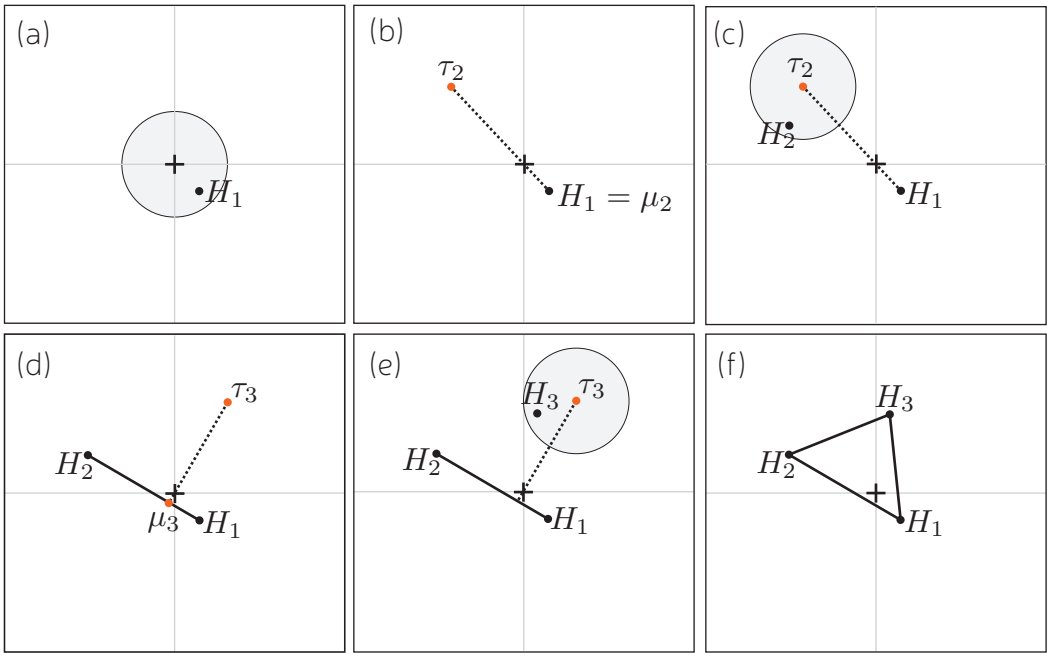

FIG. 2. The geometric intuition of the convex hull finding algorithm. The cross marks the origin corresponding to V. (a) We find a $U_1 = Ve^{iH_1}$ so that H_1 is near the origin. (b) We extrapolate from $\mu_1 = H_1$ through the origin to a point τ_2. (c) We find a $U_2 = Ve^{iH_2}$ close to $Ve^{i\tau_2}$, so that H_2 is near to τ_2. (d) We form the convex hull of H_1 and H_2 and find the point μ_3, which is closest to the origin. From here we extrapolate out through the origin to the point τ_3. (e) We find a $U_3 = Ve^{iH_3}$ close to $Ve^{i\tau_3}$, so that H_3 is near to τ_3. (f) We form the convex hull of H_1, H_2 and H_3 and find the origin lies inside the hull, and so the algorithm terminates. Note that none of the H_j can stray far from the origin.

IV. GENERAL ROTATIONS

We show here that Thm. 1 follows from Lem. 1. First, let $\mathcal{G}_\epsilon$ be the subset of $\langle\mathcal{G}\rangle$ such that they can be synthesized with cost not exceeding $f(\epsilon)$. We have that $\mathcal{G}_\epsilon$ is an ϵ-cover of $\mathcal{L}$. That is, for all $V \in \mathcal{L}$ there exists a $U \in \mathcal{G}_\epsilon$ with $||U - V|| \leq \epsilon$. Since we work with a unitarily invariant norm this can be restated as $||V^\dagger U - \mathbb{1}|| \leq \epsilon$. We shift to a Hermitian representation and define a H such that $U = Ve^{iH}$. Since $U \sim V$ we can choose H to have small norm, which we verify later. Our goal is to not just find a single U close to V but a whole set $\{U_j\}_j$ that allows us to use the following

Lemma 2 *Let $\{H_j\}_j$ be a set of bounded Hermitian operators $||H_j|| \leq c$ for all j. Assume, the origin lies within the convex hull $0 \in \mathrm{Conv}[\{H_j\}_j]$ with convex decomposition $0 = \sum_j p_j H_j$. It follows that*

1. *$||e^{iH_j} - \mathbb{1}|| \leq c + \frac{c^2}{2}$ for all j;*
2. *$||\sum_j p_j e^{iH_j} - \mathbb{1}|| \leq \frac{c^2}{2}$.*

When $U_j = Ve^{iH_j}$ for some unitary V, this can be restated as

1. *$||U_j - V|| \leq c + \frac{c^2}{2}$ for all j;*
2. *$||\sum_j p_j U_j - V|| \leq \frac{c^2}{2}$.*

Clearly, such a set of Hermitian operators would allow us to use Lem. 1 with constants related by $a = c + \frac{c^2}{2}$ and $b = \frac{c^2}{2}$, yielding an upperbound of $a^2 + 2b = O(c^2)$. The lemma is proved by expanding the exponentials into a power series and using standard norm properties, as shown in App. A.

The key point is that we seek a set of Hermitian operators, such that the origin is contained within the convex hull of these points. Next, we present an explicit method for finding such a convex decomposition of Hermitian operators. We assume access to an oracle performing the relevant gate-synthesis decompositions. We outline the algorithm for finding a suitable convex set containing the origin.

Convex hull finding algorithm

1. Call oracle to find U_1 such that $||U_1 - V|| \leq \epsilon$;
2. Find principle H_1 such that $U_1 = Ve^{iH_1}$;
3. Set $n = 2$, and while $0 \notin \mathrm{Conv}[\{H_j\}_{1\leq j\leq n-1}]$
   (a) Find $\mu_n \in \mathrm{Conv}[\{H_j\}_{1\leq j\leq n-1}]$ with minimum $||\mu_n||$;
   (b) Define $W_n = Ve^{i\tau_n}$ where $\tau_n := -r\epsilon\mu_n/||\mu_n||$;
   (c) Call oracle to find U_n such that $||U_n - W_n|| \leq \epsilon$;
   (d) Find principle H_n such that $U_n = Ve^{iH_n}$ and append to set $\{H_j\}_{1\leq j\leq n-1}$;
   (e) $n \to n + 1$ and return to start of while loop.

The algorithm has two free parameters ϵ and r (see step 3b). In our analysis we assume $\epsilon \leq 0.01$, and for all practical applications this is easily satisfied. We take $r = 2$ for simplicity, and the exact constants in our bounds and convergence rates depend on this choice. The algorithm behaves qualitatively the same for different r settings, assuming $\epsilon^{-1} \gg r > 1$. The algorithm has two important properties that we discuss below, leaving technical details until the appendices. The basic geometric intuition behind the algorithm is illustrated in Fig. 2.

First, for all H_j found by the algorithm we have

$$||H_j|| \leq 3\epsilon + 7\epsilon^2, \tag{23}$$

which we show in App. B. This provides us with the value $c = 3\epsilon + 7\epsilon^2$ to be substituted into Lem. 2, which traced back leads to the diamond norm upper bound

$$\begin{aligned} d_\diamond(\mathcal{E},\mathcal{U}) &\leq \frac{1}{2}(a^2 + 2b) = \frac{1}{2}\left[\left(c + \frac{1}{2}c^2\right)^2 + c^2\right] \\ &\leq 10\epsilon^2, \end{aligned} \tag{24}$$

where the last line uses $\epsilon < 0.01$ to simplify higher order terms. This gives the upper bound stated in Thm. 1.

The second important property of the algorithm is that it eventually terminates. Each U_n is distinct, and in particular its H_n falls outside the convex hull of previous points (see App. C for proof). If we further assume that there are a finite number of distinct points with bounded resource cost, then there are only a finite number of possible U_n for the algorithm to output. Since each is distinct, the algorithm must terminate in a finite number of steps. The additional assumption of a finite number of suitable point is very mild, at satisfied both for the Clifford+T gate set and also any gate set where all gates have non-zero cost. Furthermore, below we see that the algorithm need not terminate, but that sufficient iterations will work equally well.

A finite number of steps may still be very many, but we have evidence the converge is very fast. First we note that in a d-dimensional space, a simplex of $d+1$ will suffice to enclose a nontrivial volume. Though the algorithm is not ensured to converge in $d+1$ steps, it may often do so. Looking at Fig. 2, the analogous setup in Euclidean geometry hints that it will always find an enclosing simplex in $d+1$ iterations, though it is unclear whether this carries over to the topology induced by the operator norm. We can be more quantitative by considering the quantity $||\mu_n||$, which measures the distance from the convex hull. Recall that the convex hull finding algorithm halts when $||\mu_n|| = 0$. Further evidence of rapid convergence is that $||\mu_n||$ decreases exponentially fast. Specifically, we find there exists a $w > 0.62$ such that

$$||\mu_n|| < 3\epsilon e^{-wn}, \tag{25}$$

so the convergence toward zero is exponentially fast. Even exponentially small $||\mu_n||$ may be nonzero, but once $||\mu_n|| \ll \epsilon^2$ the preceding proofs can be adapted to account for nonzero $||\mu_n||$ with negligible influence on the upper bounds. All convergence proof details are given in App. C.

V. AXIAL ROTATIONS

We now consider a setting where the target V is an axial rotation of a single qubit. The only assumption we make about the generating gate set is that it contains Pauli Z as a free resource. Given a protocol for axial-synthesis, for all such V and any $\epsilon > 0$ there exists at least one U_1 such that $||U_1 - V|| \leq \epsilon$ and where U_1 has cost not exceeding $f_{\rm ax}(\epsilon)$ for some $f_{\rm ax}$. Recall that $f_{\rm ax}$ is polylogarithmic in ϵ^{-1}. For instance, the Ross-Selinger algorithm satisfies the worst case bound $f_{\rm ax}(\epsilon) \leq 4\log_2(\frac{1}{\epsilon})$, and $3\log_2(\frac{1}{\epsilon})$ on average. It will prove useful to consider $V^\dagger U_1$ and expand in the Pauli basis

$$V^\dagger U_1 = \alpha_{\mathbb{1}}\mathbb{1} + i\alpha_X X + i\alpha_Y Y + i\alpha_Z Z. \tag{26}$$

We say U_1 is an over-rotation if $\alpha_Z \geq 0$ and an under-rotation if $\alpha_Z < 0$. We require a second unitary U_2 such that the pair $\{U_1, U_2\}$ contains one over-rotation and one under-rotation. For the second rotation, we will use the Pauli expansion

$$V^\dagger U_2 = \beta_{\mathbb{1}}\mathbb{1} + i\beta_X X + i\beta_Y Y + i\beta_Z Z. \tag{27}$$

Gate-synthesis only ensures one unitary such that $||U_1 - V|| \leq \epsilon$, but a suitable U_2 can be found only slightly further away. Specifically, there must exist a suitable U_2 with cost below $f(\epsilon)$. To verify this, one constructs an axial rotation V' with $||V - V'|| = \epsilon$ and $||U_1 - V'|| > \epsilon$. Unitary gate synthesis must then provide a $U_2 \neq U_1$ within ϵ of V', such that $||U_2 - V|| \leq 2\epsilon$. Furthermore, within the same cost budget we can synthesize unitaries $U_3 = ZU_1Z$ and $U_4 = ZU_2Z$, with

$$\begin{aligned} V^\dagger U_3 &= \alpha_{\mathbb{1}}\mathbb{1} - i\alpha_X X - i\alpha_Y Y + i\alpha_Z Z, \\ V^\dagger U_4 &= \beta_{\mathbb{1}}\mathbb{1} - i\beta_X X - i\beta_Y Y + i\beta_Z Z. \end{aligned} \tag{28}$$

Considering the set $\{U_1, U_2, U_3, U_4\}$ it follows immediately that they satisfy condition (1) of Lem. 1 with $a = 2\epsilon$. Next, we assign them weights $\{p_j\} = \{\frac{1-q}{2}, \frac{q}{2}, \frac{1-q}{2}, \frac{q}{2}\}$ where $0 \leq q \leq 1$ will be fixed later. The linear combination is

$$\begin{aligned} \sum_j p_j V^\dagger U_j =& ((1-q)\alpha_{\mathbb{1}} + q\beta_{\mathbb{1}})\mathbb{1} \\ &+ i((1-q)\alpha_Z + q\beta_Z)Z. \end{aligned} \tag{29}$$

Subtracting the identity and taking the operator-norm squared,

$$\begin{aligned} ||\sum_j p_j V^\dagger U_j - \mathbb{1}||^2 =& ((1-q)\alpha_{\mathbb{1}} + q\beta_{\mathbb{1}} - 1)^2 \\ &+ ((1-q)\alpha_Z + q\beta_Z)^2. \end{aligned} \tag{30}$$

We now fix q to eliminate the second term. Considering the variables $\{\alpha_Z, \beta_Z\}$, one is positive (an over-rotation) and the other negative (an under-rotation), so zero sits within the convex hull of these variables and suitable q can be found. Specifically

$$q = \frac{\alpha_Z}{\alpha_Z - \beta_Z}, \tag{31}$$

satisfies $0 \leq q \leq 1$. With the second term cancelled and taking square roots we have

$$||\sum_j p_j V^\dagger U_j - \mathbb{1}|| = |q(\beta_{\mathbb{1}} - \alpha_{\mathbb{1}}) + (1 - \alpha_{\mathbb{1}})|\,. \quad (32)$$

By the triangle inequality and $|q| \leq 1$, we have

$$||\sum_j p_j V^\dagger U_j - \mathbb{1}|| \leq |\beta_{\mathbb{1}} - \alpha_{\mathbb{1}}| + |\alpha_{\mathbb{1}} - 1|. \quad (33)$$

Inserting $1 - 1 = 0$, so that $\beta_{\mathbb{1}} - \alpha_{\mathbb{1}} = (\beta_{\mathbb{1}} - 1) + (1 - \alpha_{\mathbb{1}})$, and again using the triangle inequality, we arrive at

$$||\sum_j p_j V^\dagger U_j - \mathbb{1}|| \leq |\beta_{\mathbb{1}} - 1| + 2|\alpha_{\mathbb{1}} - 1|. \quad (34)$$

From $||V^\dagger U_1 - \mathbb{1}|| \leq \epsilon$ we can infer that

$$||(\alpha_{\mathbb{1}} - 1)\mathbb{1} + i\alpha_X X + i\alpha_Y Y + i\alpha_Z Z||^2 \leq \epsilon^2. \quad (35)$$

Evaluating the left hand side, we obtain

$$(\alpha_{\mathbb{1}} - 1)^2 + \alpha_X^2 + \alpha_Y^2 + \alpha_Z^2 \leq \epsilon^2. \quad (36)$$

Unitarity of $V^\dagger U_1$ entails that $\alpha_{\mathbb{1}}^2 + \alpha_X^2 + \alpha_Y^2 + \alpha_Z^2 = 1$ and after some simplification, we find

$$\begin{aligned}(\alpha_{\mathbb{1}} - 1)^2 + \alpha_X^2 + \alpha_Y^2 + \alpha_Z^2 &= (\alpha_{\mathbb{1}} - 1)^2 + (1 - \alpha_{\mathbb{1}}^2), \\ &= 2(1 - \alpha_{\mathbb{1}}) \leq \epsilon^2. \end{aligned} \quad (37)$$

From which we infer $|1 - \alpha_{\mathbb{1}}| \leq \epsilon^2/2$. Similarly, from $||V^\dagger U_2 - \mathbb{1}|| \leq 2\epsilon$ we can infer $|1 - \beta_{\mathbb{1}}| \leq 2\epsilon^2$. Substituting into Eq. (34), we have

$$||\sum_j p_j V^\dagger U_j - \mathbb{1}|| \leq 3\epsilon^2, \quad (38)$$

Therefore, we have demonstrated both the necessary conditions of Lem. 1 with $a = 2\epsilon$ and $b = 3\epsilon^2$. Applying the Lemma, our channel satisfies

$$d_\diamond(\mathcal{E}, \mathcal{V}) \leq \frac{1}{2}(a^2 + 2b) \leq 5\epsilon^5. \quad (39)$$

A smaller factor than 5 is likely to be provable.

VI. CONCLUSIONS

We have seen that resource costs of fault-tolerant quantum computing can be reduced by switching to a randomised approach to gate-synthesis. It may seem counterintuitive that a randomisation process can be advantageous. However, convexity of the diamond distance naturally entails that mixing over channels of similar noise levels can only reduce the noise.

We presented a convex hull finding algorithm for finding the suitable mixing ratios. While this algorithm is exponentially fast, it is plausible that a constant time algorithm exists. We suspect that a variant of Delaunay triangulation could be used to quickly identify a suitable simplex. However, our literature search on Delaunay triangulation has only found results on Euclidean space and we have yet to ascertain if such tools carry over to the operator norm topology.

This work has only considered mixing over unitary channels, which prompts the question whether more general quantum channels might be useful. Probabilistic quantum circuits with fallback [8] is an approach to gate-synthesis that is not entirely unitary, though it makes use of an ancillary qubit and works very differently to the approach presented here. As remarked earlier, mixing can be useful in preparation of different magic states [16]. We ponder whether all these approaches can be understood within a single framework of quantum channel synthesis.

Acknowledgements.- We thank Yuan Su, Vadym Kliuchnikov, Steve Flammia, Jens Eisert, Mark Howard, Luke Heyfron, Neil J. Ross, Mark Pearce and Scott Vinay for related discussions. Several of these discussions were facilitated by the Centro de Ciencias de Benasque. This work was supported by the EPSRC (EP/M024261/1).

Appendix A: Convex hull proof

This section will prove Lem. 2. We start by showing another general result that we use in several places. Let M be a Hermitian operator with eigenvalues λ_k, so that by definition $|\lambda_k| \leq ||M||$ for all k. We consider the operator

$$e^{iM} - (\mathbb{1} + iM) = \sum_{n=2}^{\infty} \frac{1}{n!}(iM)^n. \quad (A1)$$

This can be diagonalised in the eigenbasis of M and has eigenvalues $f_M(\lambda_k) := e^{i\lambda_k} - 1 - i\lambda_k$. Therefore, we have

$$||e^{iM} - (\mathbb{1} + iM)|| = \max_k |f_M(\lambda_k)|. \quad (A2)$$

On the interval $|x| \leq \pi$, one can verify that $|e^{ix} - 1 - ix| \leq \frac{1}{2}x^2$, and so provided $||M|| \leq \pi$ we have

$$||e^{iM} - (\mathbb{1} + iM)|| \leq \frac{1}{2}||M||^2. \quad (A3)$$

Now turning specifically to Lem. 2, we have

$$||e^{iH_j} - \mathbb{1}|| = ||\sum_{n=1}^{\infty} \frac{1}{n!}(iH_j)^n|| \quad (A4)$$

$$\leq ||H_j|| + ||\sum_{n=2}^{\infty} \frac{1}{n!}(iH_j)^n||. \quad (A5)$$

Since we always choose the principle H_j, we have $||H_j|| \leq \pi$ and we can use Eq. A3 to find

$$||e^{iH_j} - \mathbb{1}|| \leq ||H_j|| + \frac{1}{2}||H_j||^2 \leq c + \frac{1}{2}c^2. \quad (A6)$$

Recall that in Lem. 2 we defined c so that $||H_j|| \leq c$ for all H_j, which explains the second inequality. Therefore,

$||e^{iH_j}|| \leq c + \frac{c^2}{2}$. This shows property (1) of Lem. 2. Next we consider the convex sum of unitaries,

$$\sum_j p_j e^{iH_j} = \mathbb{1} + (\sum_j i p_j H_j) + \sum_j p_j \sum_{n=2}^{\infty} \frac{(iH_j)^n}{n!}, \quad \text{(A7)}$$

which is split into zeroth, first and higher order terms. By assumption the linear terms vanish. Therefore,

$$\begin{aligned} ||\sum_j p_j e^{iH_j} - \mathbb{1}|| =& ||\sum_j p_j \sum_{n=2}^{\infty} \frac{(iH_j)^n}{n!}||, \quad \text{(A8)} \\ &\leq \sum_j p_j ||\sum_{n=2}^{\infty} \frac{(iH_j)^n}{n!}|| \\ &\leq \sum_j p_j \frac{c^2}{2} = \frac{c^2}{2}. \end{aligned}$$

Going from second to third line, we have again used Eq. A3. This proves Lem. 2.

Appendix B: Bounding $||H_n||$.

We wish to upper bound $||H_n||$ in terms of ϵ, the precision to which gate synthesis is assessed. The operator H_n is chosen so that e^{iH_n} provides a certain unitary, U_n, and the eigenvalues are chosen within the interval $[-\pi, \pi)$. Furthermore, on this interval one has that all eigenvalues θ satisfy $|\theta| \leq |e^{i\theta} - 1| + \frac{1}{2}|e^{i\theta} - 1|^2$. It follows that

$$||H_n|| \leq ||e^{iH_n} - \mathbb{1}|| + \frac{1}{2}||e^{iH_n} - \mathbb{1}||^2. \quad \text{(B1)}$$

Next, we note that for each $n > 1$ we have

$$\begin{aligned} ||e^{iH_n} - \mathbb{1}|| &= ||U_n - V|| \quad \text{(B2)} \\ &\leq ||U_n - W_n|| + ||W_n - V|| \\ &\leq \epsilon + ||e^{i\tau_n} - \mathbb{1}|| \\ &\leq \epsilon + ||\tau_n|| + \frac{||\tau_n||^2}{2} \\ &\leq 3\epsilon + 2\epsilon^2. \end{aligned}$$

The $n = 1$ case is similar but without the $||W_n - V||$ contribution. Combining this with Eq. (B1) we have

$$||H_n|| \leq \left(3\epsilon + 2\epsilon^2\right) + \frac{1}{2}\left(3\epsilon + 2\epsilon^2\right)^2. \quad \text{(B3)}$$

Assuming $\epsilon < 0.01$ this can be simplified to

$$||H_n|| \leq 3\epsilon + 7\epsilon^2, \quad \text{(B4)}$$

as reported in the main text. This gives the value of c for Lem. 2.

Appendix C: Convergence proof

Next we show that each U_n is new by showing the strictly monotonic decrease of $||\mu_n||$. Furthermore, we show exponential decrease of $||\mu_n||$ with n. We begin by translating the closeness of U_n to W_n into the space of Hermitian operators. We define

$$\Delta_n := H_n - \tau_n, \quad \text{(C1)}$$

and later will find an upper bound on $||\Delta_n||$. First we use these operators to construct a point in the new convex hull. Mixing H_n and μ_n gives a point in the convex hull, which must have norm no larger than $||\mu_{n+1}||$, so that

$$\begin{aligned} ||\mu_{n+1}|| &\leq ||\lambda H_n + (1-\lambda)\mu_n|| \quad \text{(C2)} \\ &= ||\mu_n \left[1 - \lambda\left(1 + 2\frac{\epsilon}{||\mu_n||}\right)\right] + \lambda\Delta_n||. \end{aligned}$$

If we consider when

$$\lambda = \left(1 + 2\frac{\epsilon}{||\mu_n||}\right)^{-1} = \frac{||\mu_n||}{||\mu_n|| + 2\epsilon}, \quad \text{(C3)}$$

then it is easy to see $0 < \lambda < 1$ and that the square bracket vanishes so that

$$||\mu_{n+1}|| \leq \lambda||\Delta_n|| = \frac{||\mu_n||}{||\mu_n|| + 2\epsilon}||\Delta_n||. \quad \text{(C4)}$$

This iteration begins with $||\mu_1||$. While the algorithm only explicitly uses μ_n for $n \geq 2$, it is natural to take μ_1 as a point in the convex hull of the singleton set $\{H_1\}$, and so $\mu_1 = H_1$. Further progress requires an upper bound on $||\Delta_n||$, which we now take a lengthy detour to find.

Adding several terms of the form $0 = (x - x)$ to Δ_n, we have

$$\begin{aligned} \Delta_n =& (-i\mathbb{1} + H_n + ie^{iH_n}) + (iV^\dagger W_n - ie^{iH_n}) \quad \text{(C5)} \\ &+ \left(-iV^\dagger W_n + i\mathbb{1} - \tau_n\right). \end{aligned}$$

Taking the norm and applying triangle inequality, we get

$$\begin{aligned} ||\Delta_n|| \leq& ||\mathbb{1} + iH_n - e^{iH_n}|| + ||V^\dagger W_n - e^{iH_n}|| \quad \text{(C6)} \\ &+ ||\mathbb{1} + i\tau_n - V^\dagger W_n|| \\ =& ||\mathbb{1} + iH_n - e^{iH_n}|| + ||W_n - U_n|| \\ &+ ||\mathbb{1} + i\tau_n - e^{\tau_n}||. \quad \text{(C7)} \end{aligned}$$

For the middle term we know $||W_n - U_n|| \leq \epsilon$, and for the first and last terms we again Eq. A3, so that

$$\begin{aligned} ||\Delta_n|| &\leq \frac{1}{2}||H_n||^2 + \epsilon + \frac{1}{2}||\tau_n||^2 \quad \text{(C8)} \\ &\leq \frac{1}{2}(3\epsilon + 7\epsilon^2)^2 + \epsilon + \frac{1}{2}(2\epsilon)^2. \end{aligned}$$

We can again use $\epsilon \leq 0.01$ to bound higher order terms to obtain

$$||\Delta_n|| \leq \epsilon + 7\epsilon^2. \quad \text{(C9)}$$

Plugging this in Eq. (C4), we have

$$||\mu_{n+1}|| \leq ||\mu_n|| \frac{\epsilon + 7\epsilon^2}{||\mu_n|| + 2\epsilon} \tag{C10}$$
$$< ||\mu_n|| \frac{1}{2} (1 + 7\epsilon) ,$$

where we have used that $0 < ||\mu_n||$. Iterating this argument n times we find exponential behaviour

$$||\mu_{n+1}|| < ||\mu_1|| e^{-wn}, \tag{C11}$$

where $w = \ln(2) - \ln(1 + 7\epsilon)$. Using our earlier assumption that $0 < \epsilon < 0.01$ guarantees that $0.69315 > w > 0.62548$. In most instances convergence will be much faster than ensured by this proof, often jumping to $||\mu_{n+1}|| = 0$ within only a few iterations. Last we note that $\mu_1 = H_1$ and that $||H_1|| \leq ||V - U_1|| + \frac{1}{2}||V - U_1||^2 \leq \epsilon + \frac{1}{2}\epsilon^2$, which gives

$$||\mu_n|| < \epsilon \left(1 + \frac{1}{2}\epsilon\right) e^{-w(n-1)} \tag{C12}$$
$$= \epsilon \left[\left(1 + \frac{1}{2}\epsilon\right) e^{w}\right] e^{-wn}.$$

Since $e^{-w} > 1/2$, we have $e^w < 2$. Combined with $\epsilon < 0.01$ we know the square bracket cannot exceed 3, which leads to Eq.(25).

[1] A. Y. Kitaev, A. Shen, and M. N. Vyalyi, *Classical and quantum computation*, Vol. 47 (American Mathematical Society Providence, 2002).
[2] C. M. Dawson and M. A. Nielsen, Quantum Info. Comput. **6**, 81 (2006).
[3] A. G. Fowler, Quantum Information and Computation **11**, 867 (2011).
[4] T. T. Pham, R. Van Meter, and C. Horsman, Physical Review A **87**, 052332 (2013).
[5] V. Kliuchnikov, D. Maslov, and M. Mosca, Phys. Rev. Lett. **110**, 190502 (2013).
[6] M. Amy, D. Maslov, M. Mosca, and M. Roetteler, Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on **32**, 818 (2013).
[7] N. J. Ross and P. Selinger, arXiv preprint arXiv:1403.2975 (2014).
[8] A. Bocharov, M. Roetteler, and K. M. Svore, Phys. Rev. Lett. **114**, 080502 (2015).
[9] Y. R. Sanders, J. J. Wallman, and B. C. Sanders, New Journal of Physics **18**, 012002 (2015).
[10] R. Kueng, D. M. Long, A. C. Doherty, and S. T. Flammia, Phys. Rev. Lett. **117**, 170502 (2016).
[11] E. Knill, arXiv preprint quant-ph/0404104 (2004).
[12] O. Kern, G. Alber, and D. L. Shepelyansky, The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics **32**, 153 (2005).
[13] V. Kliuchnikov, D. Maslov, and M. Mosca, arXiv preprint arXiv:1212.6964 (2012).
[14] J. J. Wallman and J. Emerson, Phys. Rev. A **94** (2016).
[15] G. C. Knee and W. J. Munro, Phys. Rev. A **91**, 052327 (2015).
[16] E. T. Campbell and J. O'Gorman, arXiv preprint arXiv:1603.04230 (2016).
[17] M. B. Hastings, "Turning gate synthesis errors into incoherent errors," ArXiv:1612.01011v1.
[18] R. A. Horn and C. R. Johnson, *Matrix Analysis* (Cambridge, 2006).
[19] R. Bhatia, *Matrix analysis*, Vol. 169 (Springer Science & Business Media, 2013).
[20] A. Y. Kitaev, Russian Mathematical Surveys **52**, 1191 (1997).
[21] D. Aharonov and M. Ben-Or, in *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing* (ACM, 1997) pp. 176–188.
[22] J. Emerson, R. Alicki, and K. Życzkowski, Journal of Optics B: Quantum and Semiclassical Optics **7**, S347 (2005).
[23] J. Emerson, M. Silva, O. Moussa, C. Ryan, M. Laforest, J. Baugh, D. G. Cory, and R. Laflamme, Science **317**, 1893 (2007).
[24] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Phys. Rev. A **77**, 012307 (2008).
[25] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Phys. Rev. A **80**, 012304 (2009).
[26] D.-S. Wang, D. W. Berry, M. C. de Oliveira, and B. C. Sanders, Phys. Rev. Lett. **111**, 130504 (2013).
[27] D.-S. Wang and B. C. Sanders, New Journal of Physics **17**, 043004 (2015).
[28] S. Bravyi and A. Kitaev, Phys. Rev. A **71**, 022316 (2005).
[29] S. Bravyi and J. Haah, Phys. Rev. A **86**, 052329 (2012).
[30] A. M. Meier, B. Eastin, and E. Knill, Quant. Inf. and Comp. **13**, 195 (2013).
[31] C. Jones, Phys. Rev. A **87**, 042305 (2013).